

Chapter 1: Introduction

Nuclei of light atoms can fuse together into heavier nuclei, with the accompanying release of up to about 0.4% of their rest mass energy, depending on the reactants and products. To adapt this process for practical power generation requires confining the very hot reactants until a large fraction of the fuel has fused. Three distinct means are known for maintaining the fuel at high density and fusion temperatures for long enough for a large proportional “burn”—gravitational confinement, magnetic confinement, and what is called “inertial confinement.”

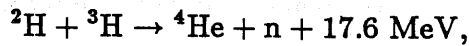
Inertial confinement fusion, or ICF, is a generalization of the hydrogen bomb mechanism. This process has been under development for application to sub-milligram fuel pellets since about 1960 [1], using lasers or beams of light charged particles rather than a fission bomb to drive the fuel to the conditions required for nuclear fusion. In addition, in 1974, Maschke [2] proposed that beams of heavy ions with parameters appropriate to pellet fusion requirements might be generated using the existing particle accelerator technology developed for high-energy particle physics research. A major extrapolation of existing technology is necessary; for this application, it is necessary to transport beams with currents in the multi-kiloampere range over long distances, while maintaining a low transverse beam temperature.

The goal of this thesis work was to determine the maximum brightness and average current density maintainable for a space-charge dominated beam in a long focusing channel. This subject is not well-treated in the conventional accelerator literature, because effects other than intrinsic space-charge dominated beam interactions with the focusing structure have limited the intensity of the beam. Circular accelerators, for example, may not have space-charge forces of more than a few percent of the average external focusing forces.

The reason is that the nonlinear space-charge force increases the spread in the transverse oscillation frequency of the particles. As a result, some particles become resonant with the unavoidable errors in the focusing channel [3,4,5]. Conventional linear accelerators (linacs) have used radio-frequency cavity acceleration and have been current-limited by, for example, the particle source current or space-charge forces in the low-velocity segment of the accelerator. For these and other reasons, alternating gradient focusing has not hitherto been explored for space-charge dominated beams. The heavy ion induction linac approach to inertial confinement fusion (ICF), however, centers about space-charge dominated transport of high-brightness beams, and the practical limitations are of overriding interest [6,7].

1.1 Approaches to Fusion

The fusion reaction requiring the least energy to initiate is the deuterium-tritium reaction



requiring a minimum temperature of several keV, and reaching the peak reaction rate (assuming Maxwellian velocity distributions for the reactants) at about 50 keV. Lawson [8] quantified the requirements for density and confinement time for thermal distributions, obtaining the result called the Lawson criterion: for the output of fusion energy from the confined fuel greater than that required to heat the fuel and overcome radiative losses, the product of the number density, n , and the confinement time (for the energy, rather than the particles), τ , must satisfy

$$n\tau \geq 10^{14} \text{ sec cm}^{-3},$$

for the reaction between deuterium and tritium, at a temperature above 10 keV. For other reactants, the temperature and $n\tau$ product required are larger.

Of the confinement mechanisms listed above, gravitational confinement is clearly unworkable on a terrestrial scale. In order to confine a fusion plasma by gravitational self-attraction, a mass of the order of the solar mass is required, and the resulting star is not a useful example.

Magnetic confinement schemes attempt to impede the free loss of fuel by applying a magnetic field within the confinement region. The presence of the magnetic field has no effect on the ultimate thermal equilibrium state of a system [9]. The purpose of the field is to slow the rate of approach to equilibrium and greatly increase the plasma confinement time. Work in this field is ongoing and has resulted in confinement times and temperatures within about an order of magnitude of what is required for "scientific breakeven," or the release of fusion energy equal to the energy used to confine and heat the plasma.

Inertial confinement derives its name from two inertial effects important in pellet fusion. A mass of fuel is accelerated radially inward to a high velocity, on the order of 2×10^7 cm/sec [6]. The inertia of the fuel carries it inward, overcoming the increasing pressure of the material and drastically increasing the density of the fuel. In pellet applications this radial acceleration is driven by heat ablation of the outer layers of the pellet. The input energy is supplied by an external driver, with the compressional force supplied by the rocket-like ablation of the heated outer material of the pellet.

The second inertial effect limits the disassembly of the dense fuel mass. The fuel particles have an average velocity determined by the temperature of the fuel. The fuel cannot escape from the high-density region in a time much shorter than that given by the ratio of the pellet radius to the thermal velocity of the ions, even in the case of a transparent plasma. However, the plasma is quite opaque to the ions, even at a temperature of 20 keV. For 1

mg of D-T compressed into a sphere 0.1 mm in radius, with a temperature of 20 keV, the 90° scattering length is on the order of 0.001 mm. Because the density is proportional to R^{-3} , where R is the minimum radius of the pellet, and the disassembly time τ is roughly proportional to R , the product of the disassembly time and the compressed fuel density is proportional to R^{-2} for a fixed fuel temperature. If the density is made sufficiently high and the temperature of the fuel is raised to about 20 keV, then the fuel simply cannot escape until a large fraction has had time to undergo fusion.

The earliest work on laser-driven ICF was directed toward the use of very small pellets, to provide a means of laboratory simulation of nuclear weapons dynamics. The initial estimates of 1–10 kJ driver energies proved to be much too low, and much of the material about the project was declassified. The possibility of applying this idea to commercial generation of electrical power was recognized, and experimental programs for both electron and light-ion beams as drivers for ICF were initiated, using the pulse power technology from flash radiography and nuclear weapons effects simulation programs. Then, in 1974, Maschke [2] proposed adapting existing particle accelerator technology to the generation of high-current, high power, heavy ion beams for use as an ICF driver (heavy ion fusion, or HIF). Recent estimates of the required pellet size and driver energy and power lie in the range of 1 mg of D-T mixture and 3–5 MJ of driver beam energy, with peak power in the hundreds of terawatts [10,11].

1.2 Heavy Ion Fusion

Hereafter we will consider only particle-beam driven ICF, and will show some advantages offered by a heavy ion beam as an ICF driver as opposed to a light particle beam. The physical properties of matter dictate certain parameters for ICF, in addition to the obvious requirement that the beam

be focused onto the pellet [6]:

- In order to obtain the required force to drive the implosion, there must be a minimum specific energy deposition (joules/gm) in the ablative portion of the pellet, on the order of 20 MJ/gm. This will raise the temperature of the ablative material to about 200 eV. The requirement provides an upper bound on the allowable range of the primary beam particles.
- The fuel compression must be nearly adiabatic to maximize the fuel density attained before incoming shock waves heat the core of the pellet and initiate thermonuclear reactions. In addition to requiring that the energy from the primary beam be deposited in the outer layer of the pellet, there must be no transport processes active to carry energy into the pellet core and increase the pressure forces opposing the implosion.
- The beam power must be high enough to deliver the energy on the time scale of the pellet implosion. To deliver megajoules of energy on the time scale of 10 ns to a pellet of surface area about 0.1 cm² requires power levels of 10¹⁴ watts and power densities of 10¹⁵ watts/cm².

These considerations constrain the range (in g/cm²) of the beam particles and thus limit the particle energy, while at the same time requiring very high beam power and total energy. For electrons or low-Z nuclei, the low energy per particle required by the range limitation necessitates very large electrical current and total charge in the beam (~ 1 MeV and 10⁸ amperes for electrons and ~ 5 MeV and 2×10^7 amperes for protons). To reduce the technical problems of transport and focusing as much as possible, Maschke [2] suggested that heavy ion beams with parameters almost within reach of current accelerator technology might be used. Beam parameters for ions

of atomic weight near 200 are expected to be 10 GeV per nucleus, with a peak beam current of 10 kiloamperes and a total charge corresponding to about 300 particle microcoulombs (μC). Although there are several unsolved problems with this approach, none presently seems prohibitive [12,13,10,11]. While these extremes of particle energy, total current, and total energy are not currently available simultaneously, beam parameters of this order are found in different accelerators around the world. The Bevalac complex at Berkeley, California, routinely accelerates such heavy nuclei as uranium to this energy, although the number of particles per pulse corresponds only about 10^{-3} – 10^{-2} particle μC . In the ATA electron induction linear accelerator at Lawrence Livermore Laboratory, beam currents of several tens of kA are attained with relativistic electrons. Beams with a total energy of several MJ are manipulated routinely within the ISR complex at CERN and at Fermilab in the USA.

Practical considerations of reaction energy deposition and wall damage [12] dictate a standoff distance of about 5–10 meters from the reactor wall and final beam lenses to the pellet. In order to focus the beam from this distance onto a pellet of radius a few mm, the optical quality of the beam must be very high. Additionally, cost and efficiency concerns require the average current density along the accelerator to be as high as possible without heating the beam transversely enough to hinder final focusing of the beam onto the pellet. This regime of near-laminar flow, space-charge dominated beam transport has until recently been treated only theoretically [14,15,16,17], with no body of experimental results to provide a comparison. How intense may a beam be before the very large electrical potential energy collectively couples to the random transverse particle kinetic energy and heats the beam? The research reported here is driven by the requirement to transport as much current as possible in each beam, subject to being able to focus the beam onto a pellet,

and to take advantage of the higher acceleration efficiency of induction linacs as the beam current becomes higher.

1.3 Early Related Work

An early study of the transport of high-current electron beams with low emittance was made by Brewer [18] in connection with traveling wave electron tube amplifiers. His experiment used solenoid focusing, and he found he could not avoid beam loss between cathode and anode unless the magnetic field strength was raised to about 1.2 times the value theoretically required for a cold beam. The cold-beam ideal case for this transport is called Brillouin flow [19], for which the current density and confining field are related (nonrelativistically) by

$$j = B_z^2 \frac{\epsilon_0 q v_z}{2m}$$

in mks units, where j is the current density, B_z is the solenoid field strength, q is the particle electric charge, and m is the particle mass.

In detailed experiments, Brewer established that his electron beam contained a significant component of thermally hot particles, which explained the need for the higher field. The particles with high transverse velocities (due to aberrations at the cathode edge) Brewer called “translaminar.” The question of the long-term stability of the transport of cold, high-current beams, particularly in an alternating gradient (A.G.) focusing array, had not been addressed before the interest in the topic for heavy ion fusion (HIF) brought it to the forefront of attention [20,17,21].

1.4 Recent results relevant to heavy ion fusion

Based on materials limits and attainable magnetic focusing strengths, Maschke [22] proposed a limit on the beam power transportable in an alter-

nating gradient channel

$$P = C \left(\frac{A}{Z} \right)^{4/3} B_Q^{2/3} (\beta_r \gamma)^{5/3} \epsilon_N^{2/3} (\gamma - 1), \quad (1.1)$$

where P is in watts, A is the atomic number of the ion, Z is its charge state in units of the electron charge, B_Q is the magnetic field strength at the pole-tips of the focusing magnets, ϵ_N is the normalized emittance (see section 2.2.1), β_r and γ are the usual relativistic factors, and C is a constant. Under certain assumptions, discussed below, Maschke showed that the value assumed by the constant is

$$C \simeq 1.67 \times 10^{15} \text{ Watt}/(\text{Tesla-meter})^{2/3}.$$

The focusing of the channel must counter both the space-charge defocusing of the beam and the spread in transverse velocities of the beam particles. The space-charge forces cancel part of the applied lens fields, leaving only part of the vacuum focusing field to contain the thermal motion of the particles. The value of the constant in the above equation is dependent on the allowable relative strengths of the channel focusing and the space-charge self-defocusing of the beam. Maschke estimated, based on his experience with radio-frequency linear accelerators, that only about half of the average restoring force provided by a focusing channel could be canceled by the space-charge self-defocusing of the beam without enabling collective interactions which would degrade the optical quality of the beam. Taking the average focusing of the channel to be linear, a single particle passing along the channel executes harmonic oscillations in the transverse dimensions, which may be characterized by a "wave number," or phase advance rate for the oscillation. This is often quoted in terms of a parameter, σ_0 , in degrees of phase advance per period of the (periodic) focusing structure (see Ch. 2).

For a given particle mass and velocity along the channel, σ_0 is proportional to the square root of the equivalent spring constant of the restoring well. For a beam with non-negligible space-charge forces, the overall restoring force is weakened, and the individual particles execute harmonic oscillations at a lower phase advance rate denoted by σ . If only a factor of two decrease in the overall focusing were possible without degrading the optical quality of the beam, then we would be constrained to have $\sigma/\sigma_0 \geq 0.7$.

Gluckstern [15] had already published a calculation for a model beam distribution (Kapchinskij-Vladimirskij, or K-V; see section 2.2.2) in the limit of a uniform focusing lattice which indicated the possibility of unstable beam behavior from purely transverse effects. This was later extended to periodic focusing structures by Smith and co-workers [16], and indicated that serious difficulties might arise if the channel focusing were raised to $\sigma_0 > 60^\circ$ (see section 2.2.1). Results published from this work by Hofmann, Laslett, Smith, and Haber [20] show many isolated patches and extended regions of instability for the K-V beam distribution.

Numerical simulation work has confirmed some of the analytical results for the K-V distribution (the only analytically known equilibrium distribution for a non-uniform focusing channel). In addition, simulation work is not restricted to this distribution, and has shown that many of the instabilities of the analytical work saturate with little or no practical effect on the beam.

The overall conclusion reached from this theoretical work was that the focusing strength of the lattice should be limited to $\sigma_0 < 60^\circ$ per focusing period in an alternating gradient channel, and the space-charge forces should not reduce the depressed phase advance, σ , below about 24° .

There have been experimental results elsewhere [23,24], as well as from this work [25], which have shown that within the limits of the lengths of the

respective channels, the space-charge forces may be so great as to almost totally cancel the external focusing. In this case, the beam temperature is so low as to be almost totally negligible, and the particle flow in the beam is nearly laminar (a particle will drift from one side of the beam to the other only after traveling through many periods of the focusing channel).

1.5 Summary of Goals and Results

For a fusion pellet driver, the emittance (see section 2.2.1) must be held below certain bounds, given by the requirement to focus the beam onto a fusion fuel pellet. The major result of this work is the measurement of the threshold for collective degradation of the emittance in high space-charge beam environments, providing information both for accelerator design in the heavy ion fusion program and for theoretical efforts to interpret the mechanism of such collective effects. To give an experimental answer to the question of beam stability we constructed a focusing channel for transport of a coasting beam (constant particle energy downstream of the beam injector). We measured the beam emittance at injection, exit, and intermediate lattice locations, to determine the low-emittance limits on the stable transport of high-current beams, resulting from collective space-charge interaction with the focusing channel. We varied the relative contributions of the space-charge and emittance terms of the envelope equations (section 2.2.2) both by attenuating the beam and by raising the transverse emittance using a set of biasable grids we could insert into the beam. We define the beam to be "stable" if the output current and emittance are equal to the values at injection.

Our results show no observable lower bound to the emittance maintainable for a given beam current (for 40 periods of quadrupole transport), as long as the focusing channel has a zero-current phase advance per period, σ_0 , less than about 90° . This result has important consequences for the design of

HIF drivers.

For $\sigma_0 > 90^\circ$, there is a band of instability in the envelope equations themselves. While we observed a definite lower bound on the emittance for $\sigma_0 > 90^\circ$ (where the envelope equations have the possibility of being unstable), the empirical limit was not well-correlated with this lowest-order instability band. Instead, the limit on beam intensity was well-described for σ_0 as high as 140° by a relation of the form of the smooth approximation relationship (derived assuming a K-V distribution for the beam)

$$\sigma^2 = \sigma_0^2 - \frac{1}{2} \left(\frac{180^\circ}{\pi} \frac{2L}{v_z} \omega_p \right)^2,$$

with a constant limiting ω_p given by

$$\omega_p \frac{2L}{v_z} \simeq \frac{2\pi}{3}.$$

Here, σ is the phase advance of the beam particles including space-charge effects, ω_p is the usual (ion) plasma frequency, v_z the beam velocity along the channel, and $2L$ is the period of the focusing channel. This value for ω_p corresponds to one "plasma oscillation" occurring while the beam transits three focusing periods.

Thus the boundary $\sigma_0 \simeq 90^\circ$ separates a region of parameter space in which the beam emittance grows very rapidly (on a length scale of ten focusing periods) from a region giving no noticeable growth ($< 10\%$) in our experimental channel of over 40 focusing periods.

Chapter 2: Beam Dynamics

2.1 Conventions

We use right-handed (x, y, z) coordinates for a linear transport system in which x is positive in the vertically upward direction and z is positive in the direction of the beam. We will use nonrelativistic formulae, relativistic corrections being negligible throughout this work. There is often confusion of the relativistic “beta”, the ratio of a velocity to the speed of light, with the “beta function” of accelerator physics, to be defined shortly. We will denote the relativistic quantity by $\beta_r = v/c$, all other occurrences of β being the accelerator function, which depends on the distance along the focusing channel. All external focusing fields will be assumed linear in transverse displacement, with a given periodicity in z . We will denote derivatives with respect to z by primes. Thus, for example, $x' \equiv \partial x / \partial z$. Finally, we will denote the RMS values of a parameter, such as the beam offset in the x dimension, by a tilde, as

$$\tilde{x} \equiv \sqrt{\langle (x - \bar{x})^2 \rangle}.$$

Space-charge will be assumed to be the only source of non-linear fields, and for most of this work only the linear part of this field is considered. We will assume a monoenergetic longitudinal beam distribution without acceleration. The experimental apparatus incorporates electrostatic quadrupoles for beam focusing, and we will write all focusing fields as electric fields, using mks units. Recall that the magnetic equivalent involves substitution of $\mathbf{y} \times \mathbf{B}$ for \mathbf{E} . The space-charge field will be calculated locally as if it were purely transverse, that is, neglecting beam envelope variations.

Periodic focusing systems are well-covered in the classic paper by Courant and Snyder [26], which includes the limitations placed on circular machines